

WE CLAIM:

1. A virtual switch for a network, the virtual switch comprising:

5 a master chassis comprising:

a first plurality of linecards; and

10 a master supervisor card for controlling the first plurality of linecards; and

a slave chassis under the control of the master supervisor card, the slave chassis comprising:

15 a second plurality of linecards; and

a slave supervisor card; and

20 a virtual switch link for communication between the master chassis and the slave chassis.

15 2. The virtual switch of claim 1, wherein the master chassis and the slave chassis communicate according to a virtual switch link protocol for logically extending a data plane of the master chassis to that of the slave chassis.

20 3. The virtual switch of claim 1, wherein the virtual switch link comprises a control virtual switch link and a data virtual switch link.

25 4. The virtual switch of claim 1, wherein the virtual switch link comprises a plurality of physical links combined to form a logical link.

30 5. The virtual switch of claim 2, wherein the virtual switch link protocol comprises a field indicating whether a packet has traversed the virtual switch link.

35 6. The virtual switch of claim 2, wherein the virtual switch link is used to synchronize routing tables of the

master chassis and the slave chassis.

7. The virtual switch of claim 3, wherein the control virtual switch link extends an internal Out-of-Band Channel to communicate between the master chassis and the slave chassis.

8. The virtual switch of claim 3, wherein the data virtual switch link extends an internal chassis data plane to communication between the master chassis and the slave chassis.

9. The virtual switch of claim 3, wherein the master supervisor communicates with the slave supervisor via inband messaging on the control virtual switch link.

10. The virtual switch of claim 3, wherein the control virtual switch link is brought on-line first and is used to determine which chassis will be the master chassis.

11. The virtual switch of claim 3, wherein a single physical link combines the control virtual switch link and the data virtual switch link.

12. The virtual switch of claim 3, wherein the control virtual switch link and the data virtual switch link are formed from separate physical links.

13. A master chassis configured to control a virtual switch for a network, the master chassis comprising:
a first plurality of linecards; and
a master supervisor card for controlling the first plurality of linecards and a slave chassis, the master supervisor card communicating with the slave chassis via a virtual switch link protocol that logically extends a

data plane of the master chassis to that of the slave chassis.

14. A method of forming a virtual switch from a plurality of physical switches in a network, the method comprising:

configuring a first physical switch as a master switch for controlling the virtual switch;

10 configuring a second physical switch as a slave switch under the control of the master switch;

forming a virtual switch link for communication between the master switch and the slave switch; and

causing the master switch and the slave switch to communicate via a virtual switch link protocol..

15 15. The method of claim 14, wherein the virtual switch link protocol comprises a source port identifier.

16. The method of claim 14, wherein the virtual switch link protocol comprises a destination port index.

20 17. The method of claim 14, wherein the virtual switch link protocol comprises source flood information.

25 18. The method of claim 14, wherein the virtual switch link protocol comprises VLAN information.

19. The method of claim 14, wherein the virtual switch link protocol indicates whether an access control list should be applied to a frame.

30 20. The method of claim 14, wherein the virtual switch link protocol indicates whether a QoS designation should be applied to a frame.

21. The method of claim 14, wherein the virtual switch link protocol indicates whether a frame is a MAC notification frame.

5 22. The method of claim 14, wherein the virtual switch link protocol includes data plane priority information for a frame.

10 23. The method of claim 14, further comprising extending a first data plane of the master switch to include a second data plane of the slave switch according to communication between the master switch and the slave switch via the virtual switch link protocol.

15 24. The method of claim 14, further comprising forming the virtual switch link from a plurality of physical links acting as a single logical link.

20 25. The method of claim 14, further comprising forming the virtual switch link to include a data virtual switch link and a control virtual switch link.

25 26. The method of claim 14, further comprising:
 updating layer 2 forwarding tables in the master chassis;

 updating layer 2 forwarding tables in the slave chassis; and

30 correcting inconsistencies between the layer 2 forwarding tables in the master chassis and the layer 2 forwarding tables in the slave chassis.

35 27. The method of claim 25, wherein step of forming the virtual switch link comprises combining the data virtual switch link and the control virtual switch link on a single physical link.

28. The method of claim 25, further comprising:
 updating layer 2 forwarding tables in the master
 chassis;

5 updating layer 2 forwarding tables in the slave
 chassis; and

 correcting inconsistencies between the layer 2
forwarding tables in the master chassis and the layer 2
forwarding tables in the slave chassis according to
10 frames transmitted on the data virtual switch link.

29. The method of claim 28, wherein the frames are MAC
notification frames.

15 30. An apparatus for forming a virtual switch from a
plurality of physical switches in a distribution layer or
a core layer of a network, the apparatus comprising:

 means for configuring a first physical switch as a
master switch for controlling the virtual switch;

20 means for configuring a second physical switch as a
slave switch under the control of the master switch;

 means for forming a virtual switch link for
communication between the master switch and the slave
switch; and

25 means for causing the master switch and the slave
switch to communicate via a virtual switch link protocol
that logically extends a data plane of the master switch
to that of the slave switch.

31. A computer program embodied in a machine-readable
30 medium, the computer program containing instructions for
controlling a plurality of physical switches of a network
to perform the following steps:

 configuring a first physical switch as a master
switch for controlling the virtual switch;

35 configuring a second physical switch as a slave

switch under the control of the master switch;
forming a virtual switch link for communication
between the master switch and the slave switch; and
causing the master switch and the slave switch to
communicate via a virtual switch link protocol.

32. The computer program of claim 31, further comprising
instructions for controlling a plurality of physical
switches of a network to logically extend a data plane of
the master switch to that of the slave switch.

10 33. A method of initializing a virtual network device,
comprising:

performing a handshake sequence between a first
chassis and a second chassis, the first chassis and the
second chassis being redundant network devices of a data
network; and

determining whether the first chassis or the second
chassis will be a master chassis for controlling a
virtual network device comprising the first chassis and
the second chassis.

20 34. The method of claim 33, wherein the handshake
sequence includes exchanging information selected from
the group consisting of a hardware version of a
supervisor; a chassis identifier; a chassis number; a
software version of each supervisor in a chassis;
hardware values for a slot in a chassis; and a slot/port
of a remote endpoint for a particular link between the
first chassis and the second chassis.

30 35. The method of claim 33, further comprising the step
of forming a control virtual switch link of the virtual
network device according to information exchanged during
the handshake sequence.

36. The method of claim 35, further comprising the step of ascertaining whether a physical link that will become a data virtual switch link is connected to both the first chassis and the second chassis.

5 37. The method of claim 36, further comprising the step of forming a data virtual switch link of the virtual network device if the ascertaining step indicated that the physical link was connected to both the first chassis and the second chassis.